

# Thermodynamic-RAM Technology Stack

M. Alexander Nugent

M. Alexander Nugent Consulting, Santa Fe, NM, USA  
KnownTech LLC, Albuquerque, NM, USA  
Xeiam LLC, Santa Fe, NM, USA  
Email: i@alexnugent.name

Timothy W. Molter

M. Alexander Nugent Consulting, Santa Fe, NM, USA  
KnownTech LLC, Albuquerque, NM, USA  
Xeiam LLC, Santa Fe, NM, USA  
Email: tim.molter@gmail.com

**Abstract**—We introduce a technology stack or specification describing the multiple levels of abstraction and specialization needed to implement a neuromorphic processor based on the theory of AHaH Computing. This specific implementation is called Thermodynamic-RAM (kT-RAM). Bringing us closer to brain-like neural computation, kT-RAM will provide a general-purpose adaptive hardware resource to existing computing platforms enabling fast and low-power machine learning capabilities that are currently hampered by the separation of memory and processing. The motivation for defining the technology stack is two-fold. First, explaining kT-RAM is much easier if it is broken down into smaller, more manageable pieces. Secondly, groups interested in realizing kT-RAM can choose a level to contribute to that matches their interest and expertise. The levels of the Thermodynamic-RAM technology stack include the memristor, Known-Synapse, AHaH Node, kT-RAM, kT-RAM instruction set, sparse spike encoding, kT-RAM emulator, and SENSE Server.

## I. INTRODUCTION

Machine learning applications span a very diverse landscape. Some areas include motor control, combinatorial search and optimization, clustering, prediction, anomaly detection, classification, regression, natural language processing, planning and inference. A common thread is that a system learns the patterns and structure of the data in its environment, builds a model, and uses that model to make predictions of subsequent events and take action. The models which emerge contain hundreds to trillions of continuously adaptive parameters. Human brains contain on the order of  $10^{15}$  adaptive synapses. How the adaptive weights are exactly implemented in an algorithm varies, and established methods include support vector machines, decision trees, artificial neural networks and deep learning, to name a few. Intuition tells us learning and modeling the environment is a valid approach in general, as the biological brain also appears to operate in this manner. The unfortunate limitation with the algorithmic approach, however, is that it runs on traditional digital hardware. In such a computer, calculations and memory updates must necessarily be performed in different physical locations, often separated by a significant distance. The power required to adapt parameters grows impractically large as the number of parameters increases owing to the tremendous energy consumed shuttling digital bits back and forth. In a biological brain (and all of Nature), the processor and

memory are the same physical substrate and computations and memory adaptations are performed in parallel. Recent progress has been made with multi-core processors and specialized parallel processing hardware like GP-GPUs, but for machine learning applications that intend to achieve the ultra-low power dissipation of biological nervous systems, it is a dead end approach. The low-power solution to machine learning occurs when the memory-processor distance goes to zero, and this can only be achieved through intrinsically adaptive hardware.

Given the success of recent advancements in machine learning algorithms combined with the hardware power dilemma, an immense pressure exists for the development neuromorphic computer hardware. The Human Brain Project and the BRAIN Initiative with funding of over EUR 1.190 billion and USD 3 billion respectively partly aim to reverse engineer the brain in order to build brain-like hardware [1], [2]. DARPA's recent SyNAPSE program funded two large American tech companies IBM and HP as well as research giant HRL labs and aimed to develop a new type of cognitive computer similar to the form and function of a mammalian brain. CogniMem is commercializing a k-nearest neighbor application specific integrated circuit (ASIC) for pattern classification, a common machine learning task found in diverse applications [3]. Stanford's Neurogrid, a computer board using mixed digital and analog computation to simulate a network, is yet another approach at neuromorphic hardware [4]. Manchester University's SpiNNaker is another hardware platform utilizing parallel cores to simulate biologically realistic spiking neural networks [5]. IBM's neurosynaptic core and TrueNorth cognitive computing system resulted from the SyNAPSE program [6]. All these platforms have yet to prove utility along the path towards mass adoption and none have solved the foundational problem of memory-process separation.

More rigorous theoretical frameworks are also being developed for the neuromorphic computing field. Recently, Traversa and Ventra have introduced the idea of 'universal memcomputing machines', a general-purpose computing machine that has the same computational power as a non-deterministic Universal Turing Machine showing intrinsic parallelization and functional polymorphism [7]. Their system and other similar proposals employ a relatively new electronic component, the memristor, whose instantaneous state is a function of its past

states. In other words, it has memory, and like a biological synapse, it can be used as a subcomponent for computation while at the same time storing a unit of data. A previous study by Thomas et al. demonstrated that the memristor can better be used to implement neuromorphic hardware than traditional CMOS electronics [8].

Our attempt to develop neuromorphic hardware takes a unique approach inspired by life, and more generally, natural self-organization. We call the theoretical result of our efforts ‘AHaH Computing’ [9]. Rather than trying to reverse engineer the brain or transfer existing machine learning algorithms to new hardware and blindly hope to end up with an elegant power efficient chip, AHaH computing was designed from the beginning with a few key constraints: (1) must result in a hardware solution where memory and computation are combined, (2) must enable most or all machine learning applications, (3) must be simple enough to build chips with existing manufacturing technology and emulated with existing computational platforms (4) must be understandable and adoptable by application developers across all manufacturing sectors. This initial motivation led us to utilize physics and biology to create a technological framework for a neuromorphic processor satisfying the above constraints.

In trying to understand how Nature computes, we stumbled upon a fundamental structure found not only in the brain but also almost everywhere one looks - a self-organizing energy-dissipating fractal that we call ‘Known’. We find it in rivers, trees, lighting and fungus, but we also find it deep within us. The air that we breathe is coupled to our blood through thousands of bifurcating flow channels that form our lungs. Our brain is coupled to our blood through thousands of bifurcating flow channels that form our arteries and veins. The neurons in our brains are built of thousands of bifurcating flow channels that form our axons and dendrites. At all scales of organization we see the same fractal built from the same simple building block: a simple structure formed of competing energy dissipation pathways. We call this building block ‘Nature’s Transistor’, as it appears to represent a foundational adaptive building block from which higher-order self-organized structures are built, much like the transistor is a building block for modern computing.

When multiple conduction pathways compete to dissipate energy through an adaptive container, the container will adapt in a particular way that leads to the maximization of energy dissipation. We call this mechanism the Anti-Hebbian and Hebbian (AHaH) plasticity rule. It is computationally universal, but perhaps more importantly and interestingly, it also leads to general-purpose solutions in machine learning. Because the AHaH rule describes a physical process, we can create efficient and dense analog AHaH synaptic circuits with memristive components. One version of these mixed signal (digital and analog) circuits forms a generic adaptive computing resource we call Thermodynamic Random Access Memory or Thermodynamic-RAM. Thermodynamics is the branch of physics that describes the temporal evolution of matter as it flows from ordered to disordered states, and

Nature’s Transistor is an energy-dissipation flow structure, hence ‘thermodynamic’.

In neural systems, the algorithm is specified by two things: the network topology and the plasticity of the interconnections or synapses. Any general-purpose neural processor must contend with the problem that hard-wired neural topology will restrict the available neural algorithms that can be run on the processor. It is also crucial that the NPU interface merge easily with modern methods of computing. A ‘Random Access Synapse’ structure satisfies these constraints.

Thermodynamic-RAM is the first attempt at realizing a working neuromorphic processor implementing the theory of AHaH Computing. While several alternative designs are feasible and may offer specific advantages over others, the first design aims to be a general computing substrate geared towards reconfigurable network topologies and the entire spectrum of the machine learning application space. In the following sections, we break down the entire design specification into various levels from ideal memristors to integrating the finished product into existing technology. Defining the individual levels of this ‘technology stack’ helps to introduce the technology step by step and group the necessary pieces into tasks with focused objectives. This allows for separate groups to specialize at one or more levels of the stack where their strengths and interests exist. Improvements at various levels can propagate throughout the whole technology ecosystem, from materials to markets, without any single participant having to bridge the whole stack. In a way, the technology stack is an industry specification.

## II. THE THERMODYNAMIC-RAM TECHNOLOGY STACK

### A. The Memristor – Metastable Switch Collection

Many memristive materials have recently been reported [10]–[14], and the trend continues. New designs and materials are being used to create a diverse range of devices. Memristor models are also being developed and incrementally improved upon [15]–[19]. Our generalized metastable switch (MSS) memristor model is to date the most accurate model shown to capture the behavior of memristors at a level of abstraction sufficient to enable efficient circuit simulations while simultaneously describing as wide a range of devices as possible [9]. A MSS is an idealized two-state element that switches probabilistically between its two states as a function of applied voltage bias and temperature. A memristor is modeled by a collection of MSSs evolving in time, which captures the memory-enabling hysteresis behavior. The MSS model can be made more complex to account for failure modes, for example by making the MSS state potentials temporally variable. Multiple MSS models with different state potentials can be combined in parallel or series to model increasingly more complex state systems.

In our semi-empirical model, the total current through the device comes from both a memory-dependent (MSS) current component,  $I_m$ , and a Schottky diode current,  $I_s$  in parallel:

$$I = \phi I_m(V, t) + (1 - \phi) I_s(V), \quad (1)$$




Fig. 1. Our generalized memristor model captures both the memory and exponential diode characteristics via metastable switch (MSS) and a Schottky diode and provides an excellent model for a wide range of memristive devices. Here we show a hysteresis plot for a Ag-chalcogenide device from Boise State University along with a fitted model.

where  $\phi \in [0, 1]$ . A value of  $\phi = 1$  represents a device that contains no Schottky diode effects. The Schottky diode effect accounts for the exponential behavior found in many devices and allows for the accurate modeling of that effect, which the MSS component cannot capture alone.

Thermodynamic-RAM is not constrained to just one particular memristive device; any memristive device can be used as long as it meets the following criteria: (1) it is incremental and (2) its state change is voltage dependent. Based on our current understanding, the ideal device would have low thresholds of adaptation ( $< 0.2$  V), on-state resistance of  $\sim 10$  k $\Omega$  or greater, high dynamic range, durability, the capability of incremental operation with very short pulse widths and long retention times of a week or more. However, even devices that deviate considerably from these parameters will be useful in more specific applications. For example, short retention times on the order of seconds are perfectly compatible with combinatorial optimizers.

We have previously shown that our generalized MSS model for memristors accurately models four potential memristor candidates [9] for Thermodynamic-RAM, and we have incorporated the model into our circuit simulation and machine learning benchmarking software. A recent Ag-chalcogenide memristor from Boise State University device and model hysteresis plot is shown in Figure 1. The model provides common ground from which the diversity of devices can be compared and incorporated into the technology stack. By modeling a device with the MSS model, a material scientist can evaluate




Fig. 2. A) A self-organizing energy-dissipating fractal that we call Known can be found throughout Nature, is composed of a simple repeating structure formed of competing energy dissipation pathways. B) We call this building-block a Known Synapse. C) A differential pair of memristors provide a means for implementing a Known synapse in our electronics. A Known synapse is like Nature's transistor.

its utility across real-world benchmarks via software emulators and gain valuable insight into which memristive properties are, and are not, useful in the application space.

### B. Known Synapse – Competing Memristors

A memristor is an adaptive energy-dissipating pathway. As current flows through it, its internal state changes and heat is exchanged to the surrounding environment. When two adaptive energy-dissipating pathways compete for conduction resources, a Known synapse (Nature's transistor) will emerge. Two competing memristors thus form a Known synapse as shown in Figure 2. We see this building block for self-organized structures throughout Nature, for example in arteries, veins, lungs, neurons, leaves, branches, roots, lightning, rivers and mycelium networks of fungus. We observe that in all cases there is a particle that flows through competitive energy-dissipating assemblies. The particle is either directly a carrier of free energy dissipation or else it appears to gate access, like a key to a lock, to free energy dissipation of the units in the collective. Some examples of these particles include

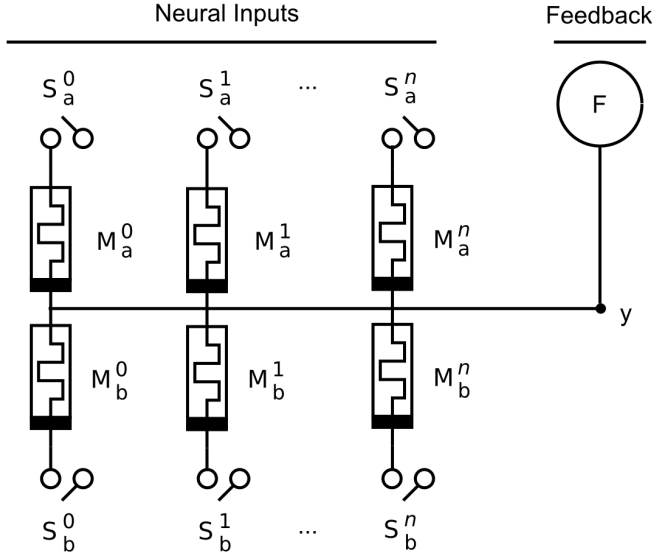


Fig. 3. An AHaH Node is made up of  $n$  Known synapses sharing a common output electrode,  $y$ . The Known synapse and the AHaH Node are analogous to a biological synapse and neuron, respectively. In Thermodynamic-RAM, the number of input synapses can be configured via software and several AHaH Nodes can be connected together to form any desired network topology by a technique called temporal partitioning.

water in plants, ATP in cells, blood in bodies, neurotrophins in brains, and money in economies. In the cases of whirlpools, hurricanes, tornadoes and convection currents we note that although the final structure does not appear to be built of competitive structures, it is the result of a competitive process with one winner; namely, the spin or rotation.

The circuits capable of achieving AHaH plasticity can be broadly categorized by the electrode configuration that forms the Known synapse as well as how the input activation (current) is converted to a feedback voltage that drives unsupervised anti-Hebbian learning [20], [21]. Synaptic currents can be converted to a feedback voltage statically (resistors or memristors), dynamically (capacitors), or actively (operational amplifiers). Each configuration requires unique circuitry to drive the electrodes so as to achieve AHaH plasticity, and multiple driving methods exist. Both polar and non-polar memristors can be used, the later requiring long periods of decay following periods of learning to prevent device saturation. The result is that a very large number of AHaH circuits exist. Herein, a ‘2-1’ two-phase circuit configuration with polar memristors is introduced because of its compactness and because it is amenable to simple mathematical analysis.

### C. AHaH Node – Collections of Known Synapses

An AHaH Node is formed when a collective of Known synapses are coupled to a common readout line. Through spike encoding and temporal multiplexing, an AHaH Node is capable of being partitioned into smaller functional AHaH Nodes. An AHaH Node provides a simple but computationally universal (and extremely useful) adaptation resource.

The functional objective of the AHaH Node shown in Figure 3 is to produce an analog output on electrode  $y$ , given an arbitrary spike input of length  $N$  with  $k$  active inputs and  $N - k$  inactive (floating) inputs. The circuit consists of one or more memristor pairs (Known synapses) sharing a common electrode labeled  $y$ . Switches gating access to a driving voltage are labeled with an  $S$ , referring to ‘spike’. The individual switches for spike inputs of the AHaH Node are labeled  $S^0, S^1 \dots S^n$ . The driving voltage source for supervised and unsupervised learning is labeled  $F$ . The subscript values  $a$  and  $b$  indicate the positive and negative dissipative pathways, respectively.

During the read phase, switches  $S_a$  and  $S_b$  are set to  $+V$  and  $-V$  respectively for all  $k$  active inputs. Inactive  $S$  inputs are left floating. The combined conductance of the active inputs produce an output voltage on electrode  $y$ . This analog signal contains useful confidence information and can be digitized via the  $\text{sgn}()$  function to either a logical 1 or a 0, if desired.

During the write phase, voltage source  $F$  is set to either  $V_y^{\text{write}} = V \text{sgn}(V_y^{\text{read}})$  (unsupervised) or  $V_y^{\text{write}} = V \text{sgn}(s)$  (supervised), where  $s$  is an externally applied teaching signal. The polarity of the driving voltage sources gates by the switches  $S$  are inverted to  $-V$  and  $+V$ . The polarity switch causes all active memristors to be driven to a less conductive state, counteracting the read phase. If this dynamic counteraction did not take place, the memristors would quickly saturate into their maximally conductive states, rendering the synapses useless.

A more intuitive explanation of the above feedback cycle is that “the winning pathway is rewarded by not getting decayed.” Each synapse can be thought of as two competing energy dissipating pathways (positive or negative evaluations) that are building structure (differential conductance). We may apply reinforcing Hebbian feedback by (1) allowing the winning pathway to dissipate more energy or (2) forcing the decay of the losing pathway. If we chose method (1) then we must at some future time ensure that we decay the conductance before device saturation is reached. If we chose method (2) then we achieve both decay and reinforcement at the same time. Method (2) is faster while method (1) is more energy efficient. The lowest energy solution is to use natural decay rather than forced decay, but this introduces complexities associated with matching the decay rate to the particular processing task.

### D. kT-RAM – AHaH Circuit with RAM Interface

As previously stated, the particular design of kT-RAM presented in this paper prioritizes flexibility and general utility above anything else, much in the same way that a CPU is designed for general purpose use. This particular design builds upon commodity RAM using its form factor and the row and column address space mapping to specific bit cells. Modifying RAM to create a kT-RAM core requires the following steps: (1) removal of the RAM reading circuitry, (2) minor design modifications of the RAM cells, (3) the addition of memristive synapses to the RAM cells, (4) addition of H-Tree circuitry connecting the synapses, (5) and addition of driving and output

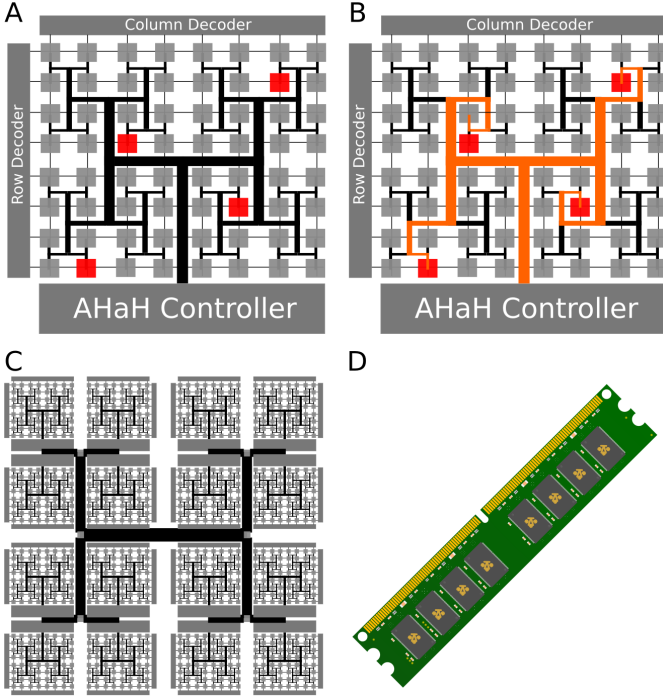


Fig. 4. A) An AHaH Circuit is superimposed on top of a normal RAM core and the synaptic inputs are turned on in the same addressable manner in which bits are set in RAM. B) During the read and write phases, the activated synapses are accessed in parallel and their individual states are concurrently adapted. C) By coupling several cores together, very large kt-RAM can be created for tasks such as inference. D) kt-RAM can borrow from existing RAM architecture to easily integrate into existing digital computing platforms.

sensing circuitry - the ‘AHaH Controller’. Multiple kt-RAM cores can be manufactured and connected to each other on the same die (Figure 4 C). Leveraging existing techniques and experience of foundries capable of producing commodity RAM as well as using three to five generation-old processing facilities will make the prototyping and manufacturing of kt-RAM relatively inexpensive. Even the final packaging of kt-RAM modules (Figure 4 D) can leverage existing commodity hardware infrastructure.

Figure 4A and B show what kt-RAM would look like with its H-Tree sensing node connecting all the underlying synapses located at each cell in the RAM array. The fractal binary tree shown is the AHaH Node’s output electrode,  $y$ , as shown in Figure 3. While at first glance it appears like this architecture leads to one giant AHaH Node per chip or core, the core can be partitioned into smaller AHaH Nodes of arbitrary size by temporally partitioning sub portions of the tree. In other words, so long as it is guaranteed that synapses assigned to a particular AHaH Node partition are never co-activated with other partitions, these ‘virtual’ AHaH Nodes can co-exist on the same physical core. This allows us to effectively exploit the extreme speed of modern electronics. Any desired network topology linking AHaH Nodes together can be achieved easily through a kt-RAM/CPU/RAM pairing. Software enforces the constraints, while the hardware remains flexible.

Through temporal partitioning combined with spike encod-

TABLE I  
kt-RAM INSTRUCTION SET

Instruction	Synapse Driving Voltage	Feedback Voltage (F)
FF	Forward-Float	None/Floating
FH	Forward-High	$-V$
FL	Forward-Low	$+V$
FU	Forward-Unsupervised	$-V$ if $y \geq 0$ else $+V$
FA	Forward-Anti-Unsupervised	$+V$ if $y \geq 0$ else $-V$
FZ	Forward-Zero	0
RF	Reverse-Float	None/Floating
RH	Reverse-High	$-V$
RL	Reverse-Low	$+V$
RU	Reverse-Unsupervised	$-V$ if $y \geq 0$ else $+V$
RA	Reverse-Anti-Unsupervised	$+V$ if $y \geq 0$ else $-V$
RZ	Reverse-Zero	0

ing, AHaH Nodes can be allocated with as few as one or as many synapses as the application requires and can be connected to create any network topology. This flexibility is possible because of a RAM interface with addressable rows and columns. Crossbar architectures, in addition to sneak-path issues, introduce a restrictive topology. While this is good for specialized applications, one cannot build a general-purpose machine learning substrate from an intrinsically restricted topology. Cores can be electrically coupled to form a larger combined core. The number of cores, and the way in which they are addressed and accessed will vary across implementations so as to be optimized for end use applications. AHaH Node sizes can therefore vary from one synapse to the size of the kt-RAM chip, while digital coupling could extend the maximal size to ‘the cloud’, limited only by the kt-Cores intrinsic adaptation rates and chip-to-chip communication.

#### E. kt-RAM Instruction Set

Thermodynamic RAM performs an analog sum of currents and adapts physically, eliminating the need to compute and write memory updates. One can theoretically exploit the kt-RAM instruction set (Table I) however they wish. However, to prevent weight saturation, one must pair ‘forward’ instructions with ‘reverse’ instructions. For example, a forward-read operation *FF* should be followed by a reverse operation (*RF*, *RH*, *RL*, *RZ*, *RA* or *RU*) and vice versa. The only way to extract state information is to leave the feedback voltage floating, and thus there are two possible read instructions: *FF* and *RF*. There is no such thing as a ‘non-destructive read’ operation in kt-RAM. Every memory access results in weight adaptation according to AHaH plasticity. By understanding how the AHaH rule works (AHaH Computing), we can exploit the weight adaptations to create, among other things, ‘self-healing hardware’. The act of accessing the information actually repairs and heals it.

#### F. Sparse Spike Encoding – Information Encoding

A spike stream is the means in which real-world data is asynchronously fed into kt-RAM. Its biological counterpart




Fig. 5. A spike-based system such as kT-RAM requires Spike Encoders (sensors), Spike Streams (wire bundles), Spike Channels (a wire), Spike Space (Number of Wires), Spike Sets or Patterns (active spike channels) and finally Spikes (the state of being active). A spike encoding is, surprisingly, nothing more than a digital code.

would be the bundles of axons of the nervous system which carry sensed information from sensing organs to and around the cortex. A sparse spike stream interface is the only option with kT-RAM, and it is used for all machine learning applications from robotic control to clustering to classification. This trait enables an application developer to leverage their knowledge and experience using kT-RAM in one domain and transfer it over to another. Spikes can directly address core synapses. The synaptic core address can thus be given by the sum of the AHaH Node's core partition index and the spike ID, which are both just integers in the spike space. Spikes enable kT-Core partitioning and multiplexing, which in turn enables arbitrary AHaH Node sizes and hence very flexible network topologies. Sparse spike encoding is also very energy and bandwidth efficient and has shown to produce state-of-the-art results on numerous benchmarks. We choose spikes because they work, and we are attempting to engineer a useful computing substrate. The fact that the spike encoding appears to match biology is of course curious, but ultimately not important to our objectives.

A collection of  $N$  synapses belong to a neuron (AHaH Node), each with an associated weight:  $\{w_0, w_1, \dots, w_N\}$ . A subset of the synapses in an AHaH Node can be activated by some input spike pattern, and the total neural activation is the voltage of the H-Tree, which can be read out on the common electrode,  $y$  by the AHaH Controller. For many input patterns,  $x$  is a sparse spiking representation, meaning that only a small subset of the spike channels are activated out of the spike space, and when they are, they are of value 1. So for a neuron with 16 inputs, one possible sparse-spike pattern would look like:  $x = \{1000001000000000\}$ . Since two of the 16 possible

inputs are active (spiking), we say that it has a sparsity of  $2/16$  or 12.5 %. Since most of the inputs are zero, we can write this spike pattern in a much more efficient way by just listing the index of the inputs that are spiking:  $x = \{0, 6\}$ .

We call  $x$  a 'spike set' or 'spike pattern' or sometimes just 'spikes'. The 'spike space' is the total number of 'spike channels', in this case 16. In some problems such as inference or text classification the spike space can get all the way up to 250,000 or more. A good way to picture it is as a big bundle of wires, where the total number of wires is the spike space and the set of wires active at any given time is the spike pattern. We call this bundle of wires and the information contained in it the 'spike stream'. The algorithms or hardware that convert data into a sparse-spiking representation are called 'spike encoders'. Your eyes, ears and nose are examples of spike encoders. A visual representation of this can be seen in Figure 5.

### G. kT-RAM Emulator – Cross-platform Universality

Thermodynamic-RAM is designed to plug into existing computing architectures easily. The envisioned hardware format is congruent with standard RAM chips and RAM modules and would plug into a motherboard in a variety of different ways. In general there are two main categories of integration. First, kT-RAM can be tightly coupled with the CPU, on the CPU die itself or connected via the north bridge. In this case, the instruction set of the CPU would have to be modified to accommodate the new capabilities of kT-RAM. Secondly, kT-RAM is loosely coupled as a peripheral device either via the PCI bus, the LPC bus, or via cables or ports to the south bridge. In these cases, no modification to the CPU's instruction set would be necessary, as the interfacing would be implemented over the generic plug-in points over the south bus. As in the case with other peripheral devices, a device driver would need to be developed. Additional integration configurations are also possible.

Given the envisioned hardware integration, kT-RAM simply becomes an additional resource that software developers have access to via an API. In the meantime, kT-RAM is implemented as an emulator running on von Neumann architecture, but the API will remain the same. Later, when the new NPU is available, it will replace the emulator, and existing programs will not need to be rewritten to benefit from the accelerated capabilities offered by the hardware. In any case, kT-RAM operates asynchronously. As new spike streams arrive, the driver in control of kT-RAM is responsible for activating the correct synapses and providing the AHaH controller with an instruction pair for each AHaH Node. The returned activation value can then be passed back to the program and used as needed.

Emulators allow developers to commence application development while remaining competitive with competing machine learning approaches. In other words, we can build a market for kT-RAM across all existing computing platforms while we simultaneously build the next generation of kT-RAM hardware. kT-RAM software emulators for both memristive



circuit validation and near-term application development on digital computers have already been developed and deployed commercially on real-world client problems. Our current digital kT-Core emulators have proven to be extremely efficient running on commodity hardware, matching and in many cases exceeding existing methods in benchmarks of solution performance, energy and memory efficiency. Thermodynamic-RAM is not a ‘ten year technology’ nor is it ‘bleeding edge’. Rather, it is already solving real-world machine learning problems on existing digital platforms.

#### H. SENSE Server – Plug-and-Play Machine Learning Apps

While a machine learning application developer using the kT-RAM Emulator would have full control of the design of the application and can use kT-RAM to its full potential, she would be required to understand the instruction set and underlying mechanics of kT-RAM and AHaH Computing. This level of development is analogous to writing assembly code or using a very low-level programming library.

To assist in the rapid development of applications based on kT-RAM, we have developed a top-level server-based framework. We call it ‘Scalable and Extensible Neural Sensing Engine’ or ‘SENSE Server’ for short. The SENSE server contains higher level pre-built machine learning modules, standard spike encoders, buffers, spike stream joiners and other miscellaneous building blocks, which can be configured by the developer for a unique machine learning application. This level of development is analogous to an SQL server like MySQL, where you provide a configuration file to specify its behavior. Like the MySQL server, the SENSE Server runs as a daemon service, waiting for asynchronous interactions from the outside world. In the case of the SENSE server, it is waiting for incoming spikes flowing in over the configured spike streams. To install and run the SENSE server on Linux, you would run a command in a terminal such as ‘sudo apt-get install knowm-sense’ followed by ‘start knowm-sense myconfig.yml’, where ‘myconfig.yml’ would be the custom configuration file defining the ‘netlist’ and parameter settings of the particular machine learning application. The SENSE server can be run on commodity computer hardware, robotic platforms or mobile devices with a Linux or \*nix-based operating system. Work is underway to port the SENSE server over to additional platforms such as iOS, Android, and Windows.

### III. CONCLUSION

In this paper, we have introduced Thermodynamic-RAM and a technology stack, a specification or blueprint, for a future industry enabled by AHaH Computing. kT-RAM is a particular design that prioritizes flexibility and general utility above anything else, much in the same way that a CPU is designed for general purpose use. The flexibility offered by this design allows for a single architecture that can be used for the entire range of machine learning applications given their unique network topologies. Much like the cortex integrates signals from different sensing organs via a common ‘protocol’, the sparse spike encoding interface of kT-RAM

allows for a well defined way to integrate environmental data asynchronously. Conveniently, the sparse spike encoding interface is a perfect bridge between digital systems and neuromorphic hardware. Just as modern computing is based on the concept of the bit and quantum computing is based on the concept of the qubit, AHaH computing is built from the ahbit. AHaH attractor states are a reflection of the underlying statistics (history) of the applied data stream. It is both the collection of physical synapses and also the structure of the information that is being processed that together result in an AHaH attractor state. Hence, an ahbit is what results when we couple information to energy dissipation. Our kT-RAM design borrows heavily from commodity RAM using its form factor to build upon and leverage today’s chip manufacturing resources. The RAM module packaging and concise instruction set will allow for easy integration into existing computing platforms such as commodity personal computers, smart phones and super computers. Our kT-RAM emulator allows us to develop applications, demonstrate utility, and justify a large investment into future chip development. When chips are available, existing applications using the emulator API will not have to be rewritten in order to take advantage of new hardware acceleration capabilities. The topmost level of the kT-RAM technology stack is the SENSE Server, a framework for configuring a custom machine learning application, based on a ‘netlist’ of pre-built machine learning modules, standard spike encoders, buffers, spike stream joiners and other miscellaneous building blocks.

### IV. FUTURE WORK

At the core of the adaptive power problem is the energy wasted during memoryprocessor communication. The ultimate solution to the problem entails finding ways to let memory configure itself, and AHaH computing is a conceptual framework for understanding how this can be accomplished. Thermodynamic-RAM is an adaptive physical hardware resource for providing AHaH plasticity and hence a substrate from which AHaH computing is possible. In previous work, we have shown demonstrations of universal logic, clustering, classification, prediction, robotic actuation and combinatorial optimization benchmarks using AHaH computing, and we have successfully mapped all these functions to the kT-RAM instruction set and emulator. Efficient emulation has already been demonstrated on commodity von Neumann hardware, and a path ahead towards neuromorphic chips has been defined. Along the way, the emulator will be ported to coprocessors like GP-GPUs, FPGAs and Epiphany chips to further improve speed and power efficiency with available hardware. Progress is being made independently at various levels, but a coordinated and focused effort by multiple participants is needed to bridge the full technology stack.

### ACKNOWLEDGMENT

The authors would like to thank the Air Force Research Labs in Rome, NY for their support under the SBIR/STTR programs AF10-BT31, AF121-049. The authors would like

to thank Kristy A. Campbell from Boise State University for graciously providing us with memristor device data.

## REFERENCES

- [1] T. Hampton, "European-led project strives to simulate the human brain," *JAMA*, vol. 311, no. 16, pp. 1598–1600, 2014.
- [2] T. R. Insel, S. C. Landis, and F. S. Collins, "The NIH brain initiative," *Science*, vol. 340, no. 6133, pp. 687–688, 2013.
- [3] B. McCormick, "Applying cognitive memory to cybersecurity," in *Network Science and Cybersecurity*. Springer, 2014, pp. 63–73.
- [4] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandrasekaran, J. Bussat, R. Alvarez-Icaza, J. Arthur, P. Merolla, and K. Boahen, "Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations," *Proceedings of the IEEE*, vol. 102, no. 5, pp. 699–716, 2014.
- [5] J. Navaridas, S. Furber, J. Garside, X. Jin, M. Khan, D. Lester, M. Luján, J. Miguel-Alonso, E. Painkras, C. Patterson *et al.*, "SpiNNaker: Fault tolerance in a power-and area-constrained large-scale neuromimetic architecture," *Parallel Computing*, vol. 39, no. 11, pp. 693–708, 2013.
- [6] S. K. Esser, A. Andreopoulos, R. Appuswamy, P. Datta, D. Barch, A. Amir, J. Arthur, A. Cassidy, M. Flickner, P. Merolla *et al.*, "Cognitive computing systems: Algorithms and applications for networks of neuromimetic cores," in *Neural Networks (IJCNN), The 2013 International Joint Conference on*. IEEE, 2013, pp. 1–10.
- [7] F. L. Traversa and M. Di Ventra, "Universal memcomputing machines," *arXiv preprint arXiv:1405.0931*, 2014.
- [8] A. Thomas, "Memristor-based neural networks," *Journal of Physics D: Applied Physics*, vol. 46, no. 9, p. 093001, 2013.
- [9] M. A. Nugent and M. T. W. "Ahah computing-from metastable switches to attractors to machine learning," *PLoS ONE*, vol. 9, p. e85175, 02 2014.
- [10] A. S. Oblea, A. Timilsina, D. Moore, and K. A. Campbell, "Silver chalcogenide based memristor devices," in *Proc. 2010 IEEE International Joint Conference on Neural Networks (IJCNN)*, 2010, pp. 1–3.
- [11] Y. Yang, P. Sheridan, and W. Lu, "Complementary resistive switching in tantalum oxide-based resistive memory devices," *Applied Physics Letters*, vol. 100, no. 20, p. 203112, 2012.
- [12] I. Valov and M. N. Kozicki, "Cation-based resistance change memory," *Journal of Physics D: Applied Physics*, vol. 46, no. 7, p. 074005, 2013.
- [13] T. Hasegawa, A. Nayak, T. Ohno, K. Terabe, T. Tsuruoka, J. K. Gimzewski, and M. Aono, "Memristive operations demonstrated by gap-type atomic switches," *Applied Physics A*, vol. 102, no. 4, pp. 811–815, 2011.
- [14] B. L. Jackson, B. Rajendran, G. S. Corrado, M. Breitwisch, G. W. Burr, R. Cheek, K. Gopalakrishnan, S. Raoux, C. T. Rettner, A. Padilla, A. G. Schrott, R. S. Shenoy, B. N. Kurdi, C. H. Lam, and D. S. Modha, "Nanoscale electronic synapses using phase change devices," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 9, no. 2, p. 12, 2013.
- [15] S. Choi, S. Ambrogio, S. Balatti, F. Nardi, and D. Ielmini, "Resistance drift model for conductive-bridge (CB) RAM by filament surface relaxation," in *Proc. 2012 IEEE 4th International Memory Workshop (IMW)*, 2012, pp. 1–4.
- [16] S. Menzel, U. Bottger, and R. Waser, "Simulation of multilevel switching in electrochemical metallization memory cells," *Journal of Applied Physics*, vol. 111, no. 1, p. 014501, 2012.
- [17] T. Chang, S.-H. Jo, K.-H. Kim, P. Sheridan, S. Gaba, and W. Lu, "Synaptic behaviors and modeling of a metal oxide memristive device," *Applied Physics A*, vol. 102, no. 4, pp. 857–863, 2011.
- [18] P. Sheridan, K.-H. Kim, S. Gaba, T. Chang, L. Chen, and W. Lu, "Device and SPICE modeling of RRAM devices," *Nanoscale*, vol. 3, no. 9, pp. 3833–3840, 2011.
- [19] D. Biolek, Z. Biolek, and V. Biolkova, "SPICE modeling of memristive, memcapacitive and meminductive systems," in *Proc. 2009 IEEE European Conference on Circuit Theory and Design (ECCTD)*, 2009, pp. 249–252.
- [20] Nugent, "Universal logic gate utilizing nanotechnology," Patent US Patent 7,420,396, 2008.
- [21] A. Nugent, "Methodology for the configuration and repair of unreliable switching elements," Patent US 7,599,895, 2009.